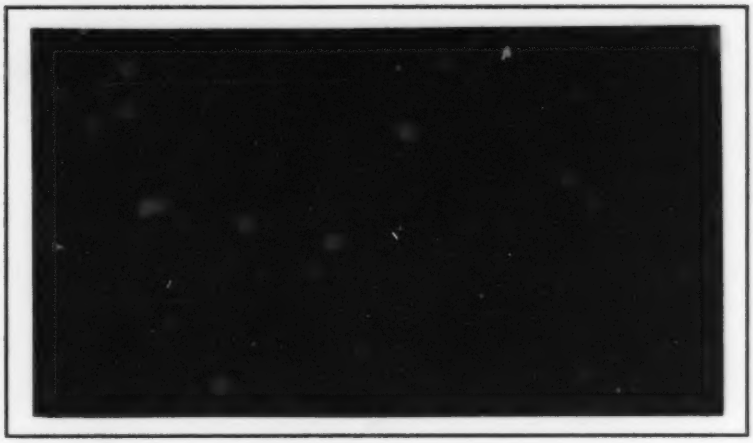

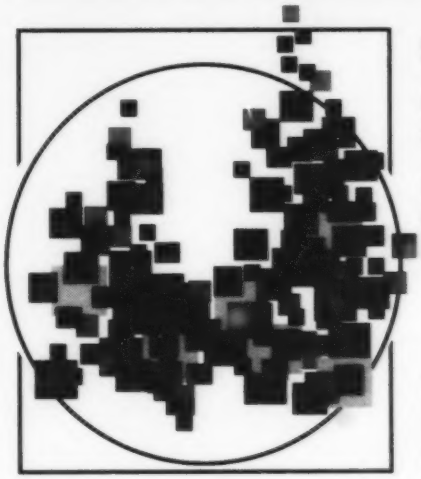

CENTRE FOR
SYSTEMS
SCIENCE

LCCR

Simon Fraser University, Burnaby, B.C., Canada V5A 1S6 / (604) 291-4704

Spectral Sharpening: Sensor Transformations for Improved Color Constancy

Graham D Finlayson, Mark S Drew and Brian V Funt

CSS/LCCR TR91-14

December 1991

Spectral Sharpening: Sensor Transformations for Improved Color Constancy

Graham D. Finlayson, Mark S. Drew and Brian V. Funt
School of Computing Science
Simon Fraser University
Vancouver, B.C.
Canada V5A 1S6
(604) 291-3458
e-mail graham@cs.sfu.ca

©1991 G.D. Finlayson, M.S. Drew and B.V. Funt

92-05825/

Abstract

We develop two sensor transformations, collectively called *spectral sharpening*, which create narrowband sensors from broadband sensors. Because changes in illumination effect narrowband sensors in simpler ways than broadband ones, spectral sharpening will improve the performance of any color constancy algorithm based on an independent adjustment of the sensor response channels. Independent adjustment of multiplicative coefficients corresponds to the application of a diagonal matrix transform (DMT) to the sensor response vector and is a common feature of many theories of color constancy: Land's retinex and von Kries adaptation in particular. Although diagonal matrix transform algorithms will achieve better color constancy with narrowband sensors, the spectral sensitivities of the human cones cannot be changed. Nonetheless, new sensors with sharper (more narrowband) spectral sensitivities can be derived as linear combinations of the cones' original spectral sensitivities. We set forth two techniques for spectral sharpening. *Sensor-based sharpening* extracts a sharpened sensor by finding the linear combination that maximizes the spectral sensitivity within a specified interval of the visible spectrum. Sensor-based sharpening focusses specifically on producing a narrowband sensor. *Model-based sharpening*, on the other hand, extracts new sensors by optimizing the ability of a DMT to account for illumination changes by examining sensor response vectors of a set of surfaces measured under different illuminants. Interestingly, model-based and sensor-based sharpening yield very similar results. When sensors based on sharpened cone sensitivities are used, a diagonal matrix transformation models illumination changes extremely well. We present experimental results suggesting that in general non-diagonal transforms can do only marginally better. Our sharpening results correlate well with the psychophysical evidence of spectral sharpening in the human visual system.

Keywords: spectral sharpening, color constancy, lightness

1 Introduction

All theories of human color constancy are strongly constrained by the spectral sensitivities of the cones. While these sensitivities obviously cannot be changed, we need not assume that they form the only possible input to the color constancy process. New sensor sensitivities can be constructed as linear combinations of the cone sensitivities. To enhance the performance of any color constancy method that responds to changing illumination by adjusting the sensor channels independently, we propose that the best linear combination to choose is the one creating the sharpest, most narrowband, sensor sensitivities.

The perceived color of an object is, more or less, independent of the illuminant under which it is observed. Thus color is a quality not of the reflected light but of the object's surface. The ability to label objects with color names, that refer only to surface reflectance properties, is called color constancy[1].

In this paper we derive two transformations of sensor sensitivities which we collectively call *spectral sharpening*. Spectral sharpening improves the performance of those algorithms that solve for color constancy by applying a diagonal matrix transform; these include Von Kries adaptation[2], Forsyth's[3] gamut-mapping approach and all the lightness algorithms. Forsyth[3] proved that, if surface reflectances are unconstrained, perfect color constancy can only be achieved via a diagonal matrix transform; this strongly motivates our study of DMT algorithms.

Spectral sharpening is a general method which can be applied to arbitrary sets of sensors. For this reason, the spectral sharpening transforms will be developed without explicit reference to the cones or limitation on the number of sensor classes.

DMT algorithms assume that the effect of the illuminant can be discounted by applying a diagonal matrix transform to each sensor response vector[1].We call these transformed vectors surface *descriptors*. If illumination conditions are constant throughout the scene, then a single diagonal transform, $\mathcal{D}$, will apply throughout the image. DMT algorithms differ only in their computational strategies for finding the diagonal matrix.

DMT color constancy is expressed mathematically in equation (1). Here $\mathbf{p}_i^e$ denotes a 3-vector of sensor responses, where i and e index surface reflectances and illumination respectively. The vector $\mathbf{p}_i$ represents a descriptor and is independent of the illuminant. Throughout the paper boldface denotes vector quantities.

$$\mathbf{p}_i \approx \mathcal{D}\mathbf{p}_i^e \tag{1}$$

In general, there may be significant error in this approximation. It appears that the narrower

[1]It is sufficient that the diagonal transform maps response vectors to their appearance under a canonical illuminant.

the chosen sensor receptivities the better the performance of DMT algorithms—indeed sensors that are sensitive to single wavelengths are known to be optimal (e.g. see [3]). These observations are the motivation for spectral sharpening.

If we can find a linear combination of sensor sensitivities such that the new sensors are sharper (more narrowband), then we predict that the performance of DMT algorithms will be improved. We consider two schemes for producing sharpened sensors: *sensor-based sharpening* and *model-based sharpening*. Both schemes yield virtually the same results and, as expected, the error associated with DMT methods using the new sensors is much lower.

Sensor-based sharpening calculates the linear combination of sensor sensitivities that maximizes the response within in a particular wavelength band. We apply sensor-based sharpening over three different ranges in order to generate three new sharpened sensors that are maximally sensitive in the long, medium, and short wave bands. Figure 1 contrasts cone fundamentals derived by Vos and Walraven[4] with a set of sharpened sensor sensitivity functions. Section 3 presents details of sensor-based sharpening.

Sensor-based sharpening does not consider the spectral properties of surfaces reflectances or illuminants. In section 4 we develop an alternate technique, called *model-based sharpening*, that sharpens sensors by examining the statistical properties of surface reflectances viewed under varying illuminants. Let V be a $3 \times n$ matrix of sensor-response vectors generated from a set S of n surfaces observed under a canonical illuminant E_1. Similarly, let W be the matrix of response vectors of S imaged under another arbitrary illuminant E_2. To the extent that DMT algorithms suffice for color constancy, V and W should be approximately equivalent under a diagonal matrix transform:

$$V \approx \mathcal{D}W \tag{2}$$

The diagonal transform is assumed to be an approximate mapping and operates within certain error bounds. These error bounds can be reduced if V and W are first transformed by a matrix $\mathcal{T}$. Model-based sharpening finds the transform $\mathcal{T}$ optimizing the diagonal matrix equivalence:

$$\mathcal{T}V \approx \mathcal{D}\mathcal{T}W$$

Transforming response vectors by $\mathcal{T}$ is equivalent to transforming the original sensor sensitivities by $\mathcal{T}$, which means that model-based sharpening also defines a set of new sensors. Our experiments show that the matrix $\mathcal{T}$ is approximately independent of the choice of illuminant E_2.

Model-based sharpening results in sensors which almost equal those obtained by sensor-based sharpening as can be seen from Figure 1. This agreement between the two very different approaches helps confirm the validity of sensor-based sharpening in general.

The results of model-based sharpening depend upon the definition of error used in the fit of the DMT. Section 5 examines the question of absolute, relative and maximal error in detail. In the light of experimental evidence, we argue that under each of the error measures, sharpened sensitivities are optimal.

Section 6 relates this work to theories of human color vision; Sharpened test and field sensitivities have been measured in humans [5, 6, 7, 8]. Foster[5], among others, advances arguments linking sharpening to color opponency. We hypothesize instead that sharpened sensor sensitivities are a natural consequence of optimizing the visual system's color constancy abilities via an initial linear transformation of the cone outputs.

2 Error Bounds and DMT Algorithms

Here we set out the mathematical model of the human visual system on which our experimental results are based. The relationship between measured spectra and sensor response vectors is given by equation (3).

$$\mathbf{p} = \int_{\omega} E(\lambda)S(\lambda)\mathbf{R}(\lambda)d\lambda \tag{3}$$

$E(\lambda)$, $S(\lambda)$, $\mathbf{R}(\lambda)$ denote illumination, surface reflectance and sensor sensitivities, respectively, and the integral is taken over the visible spectrum ω.

We consider two sets of sensor sensitivities: cone absorptance functions measured by Bowmaker[9] (BOW) and cone fundamentals derived by Vos and Walraven[4] (VW). The latter sensitivities take into account the spectral absorptions of the eye's lens and its macular pigment. Spectral sharpening improves DMT algorithm performance for both the VW and BOW sensors . Further the results for the two sensor sets are very similar. For this reason we include figures only for the VW sensors.

We use a set of 6 illuminants: 5 Judd daylight spectra [10] and CIE standard illuminant A [4]. These illuminants are applied to a set of 66 surface reflectances consisting of 12 ceramic tiles [11], the 24 Macbeth color checker [12] patches and a random selection of 30 of the natural surfaces measured by Krinov[13]. Since the Krinov, ceramic and macbeth spectra are measured in different units, all surface reflectances were normalized—they are scaled such that their squared area is equal to one. This normalization has little effect on the results presented in this paper.

We consider color constancy to be achieved if response vectors are mapped to their appearance under a canonical illuminant [3]. This is a reasonable assumption since there is no known natural illuminant $E(\lambda)$ which is constant across the visible spectrum and so the surface reflectances themselves should not be used to form descriptors. Without loss of generality we choose Judd's D55 (55

stands for 5500K) as the canonical illuminant. This implies that descriptors are response vectors for surfaces viewed under D55.

Let V and W of equation (2) be 3×66 matrices of sensor responses corresponding to the 66 surfaces viewed under two illuminants. Equation (2) can be solved for the best diagonal transformation in the least-squares sense. Since $\mathcal{D}$ is a diagonal matrix, each row of sensor responses is "fitted" independently. The components of $\mathcal{D}$ are derived as follows (note that t is the transpose operation):

$$\mathcal{D}_{ii} = V_i W_i^t [W_i W_i^t]^{-1} \tag{4}$$

Given a fixed set of sensor functions, equation (4) yields the best diagonal transformation that takes observations under one illuminant into observations under another. At this point, we are not sharpening sensors; rather, we are simply using a set of responses generated by a fixed set of sensors to find the best diagonal matrix $\mathcal{D}$ for that set of sensors. Equation (4) holds regardless of the sensors used. Later we show how to change the sensors such that the $\mathcal{D}$ found by equation (4) generates smaller errors between the transformed response vectors and their descriptors.

For both the VW and BOW sensor sensitivities, we generate experimental data to test how well DMT algorithms can possibly perform. For each illuminant, we find the optimal diagonal transform (by solving equation 4) mapping the sensor response vectors for the surface reflectance set to their appearance under the canonical illuminant D55.

Since sensor responses are 3-vectors any metric for evaluating color constancy should compare fitted vectors with their corresponding descriptors. The euclidean distance between fitted vector and descriptor, normalized with respect to the descriptor length, is a reasonable choice of metric and we will denote this normalized fitted distance as NFD.

As an example, for the VW sensitivities we used equation (4) to transform response vectors generated under D100 onto to their corresponding descriptors. Figure 2 shows the cumulative histogram of NFD error between transformed vectors and descriptors. By the time we reach an NFD of 10%, all 66 of the samples are fitted.

2.1 A Note on DMT Algorithms and Relative Error

The majority of DMT algorithms (Land[14], Horn[15] and Blake[16]) assume color ratios are independent of the illuminant. This assumption is true only if a diagonal matrix transform is an exact model for color constancy. If the percentage (relative) error of a fitted response (in one sensor channel) lies in the interval $[-\varepsilon, +\varepsilon]$ $(0 < \varepsilon < 1)$ then the corresponding percentage errors for ratios are in the range $[-\frac{2\varepsilon}{1+\varepsilon}, +\frac{2\varepsilon}{1-\varepsilon}]$. For example if the fitted percentage error is $\approx 5\%$, then ratios can

	Fitted (Relative) Error (BOW)			*Fitted (Relative) Error (VW)*		
	max red	max green	max blue	max red	max green	max blue
CIE A	33.1	30.2	21.1	25.3	18.6	21.6
D48	7.4	6.3	4.3	4.0	3.2	3.7
D65	8.7	7.0	3.9	3.9	3.76	3.5
D75	15.3	12.3	6.4	6.4	6.7	5.7
D100	26.6	20.8	9.6	10.1	11.8	8.8

Table 1: Maximum percentage relative error in each channel for diagonal transform of BOW and VW sensor responses under 5 lights to D55 , for 66 surfaces

have an error of $\approx 10\%$. Thus it is important to consider the performance of sharpened sensors with respect to relative error.

Many computational strategies for determining the diagonal matrix assign uneven weights to distinct response vectors or distinct ratios of vectors. For example, all Land's retinex methods[17, 14] will perform badly in the presence of high maximum relative error. In contrast, the gamut-mapping approach advocated by Forsyth[3] treats each distinct response vector equally and is unaffected by high maximum error. The maximum fitted relative errors for both BOW and VW sensitivities, where the diagonal matrix is computed via equation (4), are shown in Table 1.

3 Sensor-Based Sharpening

Our goal is to transform linearly the initial sensor sensitivities and create a new sharpened sensor. Thus we wish to design a function $\Psi(\mathbf{R}(\lambda), \lambda_1, \lambda_2)$ which returns the linear combination of the sensitivities $\mathbf{R}(\lambda)$ that is maximally sensitive in the closed interval $[\lambda_1, \lambda_2]$. We can solve for Ψ by minimizing the following integral:

$$\mathcal{I} = \int_{\phi} [\mathbf{R}(\lambda)^t \mathbf{c}]^2 d\lambda + \mu \left(\int_{\omega} [\mathbf{R}(\lambda)^t \mathbf{c}]^2 d\lambda - 1 \right) \tag{5}$$

where ω is the visible spectrum, ϕ denotes wavelengths outside the sharpening interval, $\mathbf{c}$ is the unknown coefficient vector, and μ is a Lagrange multiplier. The Lagrange term constrains the minimization such that the squared area under the sharpened sensor is equal to 1, which is important as the sharpened sensor can have negative sensitivities.

By differentiating with respect to $\mathbf{c}$ and equating to the zero vector, we find the stationary values of equation (5):

$$\frac{1}{2}\frac{\partial \mathcal{I}}{\partial \mathbf{c}} = \int_{\phi} \mathbf{R}(\lambda)\mathbf{R}(\lambda)^t \mathbf{c} d\lambda + \mu(\int_{\omega} \mathbf{R}(\lambda)\mathbf{R}(\lambda)^t \mathbf{c} d\lambda) = \mathbf{0} \tag{6}$$

Define the $s \times s$ (there are s initial sensors) matrix $\Lambda(\alpha) = \int_{\alpha} \mathbf{R}(\lambda)\mathbf{R}(\lambda)^t d\lambda$ so that equation (6) becomes:

$$\Lambda(\phi)\mathbf{c} + \mu\Lambda(\omega)\mathbf{c} = \mathbf{0} \tag{7}$$

Rearranging equation (7) we see that solving for $\mathbf{c}$ (and consequently the sharpened sensor) is an eigenvector problem.

$$[\Lambda(\omega)]^{-1}\Lambda(\phi)\mathbf{c} = -\mu\mathbf{c} \tag{8}$$

There may be up to s solutions of equation (8), each solution corresponds to a stationary value; Ψ returns the eigenvector which minimizes $\mathcal{I}$. It is important that $\mathbf{c}$ is a real-valued vector as this implies our sharpened sensor is a real-valued function. Indeed this is the case since both the matrices $[\Lambda(\omega)]^{-1}$ and $\Lambda(\phi)$ are positive definite and the eigenvalues of the product of two positive definite matrices are real and non-negative[18].

We sharpened the BOW functions with respect to the wavelength intervals (values are in nanometres) $[400, 480]$, $[510, 550]$ and $[580, 650]$ and the VW in the intervals $[400, 480]$, $[520, 560]$ and $[580, 650]$ [2]. These intervals were chosen to ensure that the whole visible spectrum would be sampled and that the peak sensitivities of the resulting sensors would roughly correlate with those of the cones. Sharpening in these intervals also generates sensors in close agreement to those generated by model-based sharpening, discussed in the next section. The results for the VW sensors are presented in Figure 1. Note that the sharpened curves contain negative sensitivities. These are not necessarily a cause for concern in that they do not represent negative physical sensitivities; rather they are a result of negative coefficients in a computational mechanism for color constancy.

For both the BOW and VW sensor sets, the degree of sharpening is quite significant. The peak sensitivities of the new sensors are shifted with respect to the initial sensitivities. This is due both to the choice of sharpening intervals and the shape of the BOW and VW sensitivities. The sharpened long wave mechanisms are pushed further to the long wave end of the spectrum. In contrast the medium wave mechanisms are shifted towards the shorter wavelengths. The short wave mechanisms appear unaffected by sharpening. Intriguingly, field sensitivities of the human eye measured under white light conditioning[6] with long test flashes are sharpened in analogous manner. In section 6 we review, in more detail, the psychophysical evidence for spectral sharpening.

[2] Because of the limited range of Krinov's measurements, spectra used in this paper are in the range 400nm to 650nm measured at 10nm intervals.

sensors	*% Squared Area*		
	[400, 480]	[510, 550]	[580, 650]
BOW	98.9	51.4	30.1
BOW sharpened	99.4	66.3	76.2
	[400, 480]	[520, 560]	[580, 650]
VW	97.5	67.6	40.3
VW sharpened	97.8	74.7	89.1

Table 2: Percentage of total squared area in the sharpening intervals

	Fitted relative error (BOW)			*Fitted relative error (VW)*		
	max red	max green	max blue	max red	max green	max blue
CIE A	159.7	115.5	41.4	82.5	109.2	36.4
D48	34.0	12.5	3.6	15.8	12.9	3.5
D65	42.34	11.1	2.7	19.0	11.3	2.8
D75	77.9	17.9	4.5	34.1	18.1	4.3
D100	145.6	27.3	6.9	62.0	27.1	6.1

Table 3: Maximum percentage relative errors for sharpened sensors

Table 2 contrasts the percentage squared area, lying in the sharpening intervals, for the original versus the sharpened curves. At first glance it appears that the VW sensitivities are largely unaffected by sharpening in the green region of the spectrum—the percentage squared area is increased from 67.6% to 74.7%. However, from Figure 1 it is clear that the sharpening effect is not limited to the sharpened interval.

How well can DMT algorithms perform with the sharpened BOW and VW sensors? Cumulative NFD histograms for VW sensors are plotted in Figure 3. The maximum relative errors are tabulated in Table 3. On average the fitted distance has been reduced significantly—the NFD cumulative histograms for the sharpened sensors are above those for the unsharpened ones. However, the maximum relative error has increased by a large amount. The increasing maximum can be attributed to the negative sensitivities of the new sensors. These negative regions tend to push sensor responses closer to zero and, as we consider in section 5, the relative error can dominate when response values are small.

4 Model-Based Sharpening

The statistical properties of surfaces and illuminants are not taken into account in sensor-based sharpening. It might be the case that the best sensors for DMT algorithms vary as the illumination changes. In this section we develop a model-based approach for deriving sharpened sensitivities—we optimize the sensors for DMT algorithms by examining properties of sensor response vectors.

Let S be a set of n surfaces and E_1 be the canonical illuminant and consider any E_2 such that $E_1 \neq E_2$. The $s \times n$ matrix V (there are s initial sensors) contains sensor responses for each surface viewed under E_1; each column corresponds to a single surface. Similarly, W is the $s \times n$ matrix of sensors responses for the surfaces viewed under E_2. We wish to find $s \times s$ matrices $\mathcal{T}$ and $\mathcal{D}$ ($\mathcal{D}$ is diagonal) such that $\mathcal{T}V \approx \mathcal{D}\mathcal{T}W$ is the best possible least squares approximation:

$$minimize \sum_{i,j} ([\mathcal{T}V]_{ij} - [\mathcal{D}\mathcal{T}W]_{ij})^2 \tag{9}$$

The solution for $\mathcal{D}$ is given by the Moore-Penrose inverse $\mathcal{D} = \mathcal{T}V(\mathcal{T}W)^+$, where $A^+ = A^t[AA^t]^{-1}$. After some matrix manipulation we arrive at the following condition on $\mathcal{T}$ and $\mathcal{D}$:

$$VW^t[WW^t]^{-1} = \mathcal{T}^{-1}\mathcal{D}\mathcal{T} \tag{10}$$

The matrix on the left hand side of equation (10) is similar to a unique diagonal matrix which can be calculated via the eigenvector decomposition: $VW^t[WW^t]^{-1} = \mathcal{U}\mathcal{D}\mathcal{U}^{-1}$. This implies $\mathcal{T} = \mathcal{U}^{-1}$. If we had simply asked for the best least-squares fit for some matrix $\mathcal{G}$ satisfying $V \approx \mathcal{G}W$, we would have had to solve a very similar equation, *viz.* $VW^t[WW^t]^{-1} = \mathcal{G}$. This is the non-diagonal version of equation (4). What equation (10) says, therefore, is: solve for the non-diagonal $\mathcal{G}$ and decompose it into its eigenvectors and associated diagonal matrix of eigenvalues. Therefore, $\mathcal{T}$ amounts to a change of basis transforming the initial sensors into new ones. The new sensors are such that the best least-squares solution relating $\mathcal{T}V$ and $\mathcal{T}W$ is precisely the diagonal transformation $\mathcal{D}$.

As with sensor-based sharpening, it is important that the column vectors of $\mathcal{T}$ (eigenvectors of $\mathcal{G}$) be real-valued. Depending on the properties of the data, this need not be the case. Nonetheless, the sharpened BOW and sharpened VW derived via model-based sharpening are almost identical to those derived by sensor-based sharpening (see Figure 1 for VW sensors) and we know these must be real-valued.

Results presented in the rest of the section refer to the VW sensor set. For each illuminant E_2 (not including D55), we derive the sensor transformation optimizing the performance of DMT algorithms by solving for $\mathcal{T}$ in equation (10). It turns out that the new sensors are virtually

independent of the choice of illuminant; Figure 4 show the range of the sharpened curves for the VW sensors generated under all illuminants. If we exclude CIE A from the choice of illuminants then this sharpening range diminishes, see Figure 5.

Cumulative NFD histograms for the model-based Sharpening are essentially the same as for sensor-based sharpening, shown in Figure 3, since the two sets of sensors are almost identical. Indeed such is the similarity between the sensor-based and model-based sharpened sensor that we assume they are equivalent. The results in the rest of this section refer to sharpened sensors derived via sensor-based sharpening.

Figure 6 shows the NFD histograms for the (non-diagonal) transform $\mathcal{G}'$ calculated with respect to the sharpened sensors, i.e $\mathcal{G}' = \mathcal{T}^{-1}\mathcal{D}\mathcal{T}$ where $\mathcal{T}$ is **fixed**. Note it is important that $\mathcal{T}$ does not vary with the illuminant; this is demonstrated in Figures 4 and 5. Thus we find the best diagonal transform relative to the sharpened sensors; thereafter, we assess the fit by transforming the data back to the unsharpened sensors. We do **not** at any time calculate a non-diagonal least-squares fit. It appears that $\mathcal{G}'$ performs very well in terms of the NFD statistic. This suggests that the question of sensor performance is linked to the axes with which the color space is described.

Finally in Figure 7 we contrast the cumulative NFD histograms for the non-diagonal least squares fit with those obtained for the sharpened sensors. Note the non-diagonal least-squares fit is the best that can be achieved assuming a linear model for color constancy. It appears that, excluding CIE A, a diagonal transform can achieve almost the same performance as a non-diagonal transform. Judd[19] discusses the performance of human color constancy under changing illumination. He writes that given "moderate departures from daylight in the spectral distribution of energy in the illuminant ...objects are seen ...nearly in their natural daylight colors". CIE A is a red-colored illuminant with correlated color temperature 2856K; this is cooler than all daylight illuminants and may provide an explanation of the poorer performance of constancy under CIE A.

The performance of DMT algorithms increases quite significantly when sharpened sensitivities are employed. Furthermore performance compares favorably with that of the best non-diagonal transform. If a linear transform is the central mechanism of human color constancy then our results suggest there is little to gain in calculating anything more than a diagonal transform. These results strengthen DMT theories of color constancy.

5 Other Error Measures

The experimental data presented in sections 3 and 4 strongly implies the optimality of spectral sharpening. However, sensor performance has been calculated in terms of a specific error criterion: namely, that of minimizing the sum of squared differences in each sensor channel (equation 9). As implied in section 2.1, some color constancy algorithms suggest different error metrics.

5.1 Absolute vs Relative Error

Let us consider the 3-vectors $\mathbf{p}$ and $\mathbf{q}$ where $\mathbf{p} \approx \mathbf{q}$. There are, in general, two ways to determine how closely $\mathbf{p}$ and $\mathbf{q}$ match. Absolute error is concerned with the *distance* between $\mathbf{p}$ and $\mathbf{q}$. Suitable distance measures (or metrics) include $\sum_{i=1}^{3} |\mathbf{p}_i - \mathbf{q}_i|$ and $\sum_{i=1}^{3} (\mathbf{p}_i - \mathbf{q}_i)^2$. The absolute error is often normalized by dividing by the magnitude of $\mathbf{p}$ or $\mathbf{q}$. Model-based sharpening finds sensors which are optimal in regard to minimizing the sum of absolute errors. Thus the normalized fitted distance (NFD) is a suitable performance measure.

Relative error compares the ratios of $\mathbf{p}$ and $\mathbf{q}$ to the unit vector (vector components are divided). Suppose sensor responses lie in the range $[0, 100]$; then the absolute error between responses 0.5 and 0.2 is small. However, the relative error is very large—more than 100%. This accounts for the high errors shown in Tables 1 and 3. Relative error between fitted variables can be large only if the variables have small values. Although model-based sharpening optimizes absolute error, we would predict, therefore, that on average it will perform well with respect to relative error since most sensor values will not be close to zero.

There is an important connection between relative error and the error in color ratios which allows us to evaluate the effect of sensor sharpening on the performance of DMT algorithms which employ color ratios. The maximum relative error shown in Tables 1 and 3 refers to single surface reflectances and corresponds to the ratio of a fitted 3-vector with its corresponding descriptor. Let s'_1 and s'_2 denote responses in a single sensor channel corresponding to two surfaces viewed under some illuminant and s_1 and s_2 the responses under the canonical illuminant. Let α be the coefficient which best maps s'_1, s'_2 to s_1, s_2:

$$\alpha s'_1 \approx s_1 \quad , \quad \alpha s'_2 \approx s_2$$

Without loss of generality assume $\alpha s'_1 = (1+\varepsilon_1)s_1$ and $\alpha s'_2 = (1-\varepsilon_2)s_2$ where $-1 < \varepsilon_2 < 1$. We can write the ratio $\frac{s'_1}{s'_2}$ in terms of s_1 and s_2 and thereby make a statement about ratio constancy relative to $\frac{s_1}{s_2}$:

$$\frac{s_1}{s_2} \approx \frac{(1+\varepsilon_1)s_1}{(1-\varepsilon_2)s_2}$$

It is clear that this ratio has a higher relative error than $\alpha s'_1$. In the above analysis we can subtitute ε_1 with the maximum fitted relative errors of Tables 1 and 3. Since, in all our experiments, there was always a fitted response whose error was between -1 and 1, we can conclude that there will be a color ratio with a relative error greater than or equal to the maximum fitted error.

5.2 Spectral Sharpening and Relative Error

Given a set of sensors, we wish to evaluate whether they represent a good choice in terms of relative error. Let us consider two $3 \times n$ matrices of response vectors V and W. What transform Q minimizes $V \approx QW$ in terms of relative error? This minimization is expressed in equation (11), where k indexes the sensor channels.

$$minimize\ S = \sum_{i=1}^{n}[\sum_{j=1}^{3}(\frac{Q_{kj}W_{ji}}{V_{ki}}) - 1]^2 \tag{11}$$

Keeping k fixed and differentiating with respect to Q_{kx} yields:

$$\frac{1}{2}\frac{\partial S}{\partial Q_{kx}} = \sum_{i=1}^{n}[\sum_{j=1}^{3}(\frac{Q_{kj}W_{ji}}{V_{ki}}) - 1]\frac{W_{xi}}{V_{ki}} = 0 \tag{12}$$

$$\sum_{i=1}^{n}\sum_{j=1}^{3}\frac{Q_{kj}W_{ji}W_{xi}}{V_{ki}^2} = \sum_{i=1}^{n}\frac{W_{xi}}{V_{ki}} \tag{13}$$

Let N be a $3 \times n$ matrix where $N_{ij} = \frac{W_{ji}}{V_{kj}^2}$. Similarly let M represent the $1 \times n$ row vector where $M_{1j} = \frac{1}{V_{kj}}$. Thus we can remove the explicit summations of equation (13) and rewrite it terms of matrix multiplications. The kth row of Q is denoted Q_k.

$$Q_k N W^t = MW^t \tag{14}$$

Note this is similar to, but not the same as, the best least-squares fit to $V \approx QW$ in terms of absolute error. Given the canonical illuminant E_1, some other illuminant E_2 and a set of sensors, we can generate matrices of sensor response vectors. By solving equation (14) we can determine the transform, Q mapping response vectors of E_2 to their appearance under E_1 that is optimal with respect to minimizing the total relative error. If Q is always diagonal then we conclude that the sensors are a good choice with respect to relative error. Put another way, given fixed, sharpened sensitivities can we do better than to apply a diagonal transform?

In Table 4 we tabulate the largest off-diagonal terms of the matrices Q calculated for sharpened sensors (generated by sensor-based sharpening) under different illuminants (where the rows of Q are first normalized to unit length). Note that the off-diagonal components are significant only for CIE A. The sharpened sensors perform well with respect to relative error so, on average, color ratios will be more independent of the illuminant when they are used.

	sharpened BOW	*sharpened VW*
	max	max
CIE A	0.282	0.081
D48	0.015	0.011
D65	0.019	0.015
D75	0.032	0.028
D100	0.054	0.049

Table 4: Maximum off-diagonal components

Brainard[20] presents a version of retinex where the ratio between every pair of image pixels is calculated. All ratios are treated equally and hence a high error in a single ratio will not affect the algorithm's result. Sharpened sensors, which are optimal with respect to minimizing the sum of absolute error, will improve its performance, therefore.

5.3 Minimizing the Maximum Relative Error

It is important to minimize the maximum of all relative errors or in other words, the worst case, fitted relative error. The maximum relative errors for the VOS and BOW curves are given in Table 1 and those for the corresponding sharpened sensitivities are given in Table 3. It is clear that the sharpened sensitivities are not optimal with respect to this error criterion. Land's white-patch retinex method [17] implicitly requires that the worst case error be small—a single color ratio with high relative error can significantly degrade constancy performance.

The relative error dominates only when response values are small. The negative sensitivities of the sharpened sensors increase the likelihood of small responses, which accounts for the increased maximum fitted error. Suppose we flatten (slightly) the sharpened curves and remove the negative components; then perhaps, the flatter though still sharp sensors, will perform better in terms of the worst case error. Moreover, if the degree of flattening is kept small, these curves should still perform well in terms of minimizing the sum of absolute errors.

Flattening corresponds to applying the sharpening function Ψ to wider intervals. We sharpened the BOW and VW sensor curves in the intervals [530,650] and [490,590]; the blue mechanism is unsharpened. The maximum relative errors for these sensors are shown in Table 5. As expected, in terms of maximum relative error, these sensors perform better than both the original sensors and the sharpened sensors which are optimal with respect to absolute error. The new VW sensors are graphed in Figure 8, from which it is clear that the negative components have been substantially reduced. Since the slightly flatter, sharpened curves reduce the maximum relative error of the DMT fit in equation (4), the performance of algorithms such as Land's retinex which depend on

	Fitted relative error (BOW)			*Fitted relative error (VW)*		
	max red	max green	max blue	max red	max green	max blue
CIE A	16.4	21.1	21.1	14.4	17.2	21.6
D48	3	2.7	4.3	2.7	2.2	3.7
D65	2.8	2.9	3.9	2.6	1.9	3.5
D75	4.7	5.0	6.4	4.3	3.12	5.7
D100	7.2	8.6	9.6	6.7	5.2	8.8

Table 5: Maximum relative errors when sharpening on broader intervals

sensor-response ratios should improve.

6 Spectral Sharpening and the Human Visual System

If humans employ a diagonal matrix transform for color constancy, our results show that we should expect the visual system to use sharpened sensors since doing so would optimize its performance. In this section we examine some of the psychophysical evidence for sharpened sensitivities and consider possible implications of spectral sharpening with respect to double-opponent cells.

6.1 Psychophysical Evidence

Sharpened sensitivities have been detected both in field- and test-sensitivity experiments (for a review of these terms see Foster [21]). Sperling and Hareworth[6] measured the test spectral sensitivities of humans where the eye was first conditioned to a large white background and found sharpened peaks at 530nm and 610nm and no sharpening of the blue mechanism. Our results in sharpening the VW fundamentals are consistent with these findings. They discuss a linear model of cone interactions to account for the sharpening. Specifically, they propose that the sharpened red sensor is due to the red cone minus a fraction of the green sensitivity. The sharpened green is due to the green cone minus a fraction of the red. This corresponds well with our theoretical results in that the linear combination producing the sharpened sensors involves differences in the red and green with only a slight blue contribution. The actual formulas describing our sharpened sensors are:

$$R' = 2.46R - 1.97G + 0.075B \tag{14}$$

$$G' = 1.58G - 0.66R - 0.12B \tag{15}$$

More recently, Foster[5] observed that field- and test-sensitivity spectra show sharpened peaks when derived in the presence of a small monochromatic auxiliary field coincident with the test field. Foster[22] extended this work by performing a hybrid experiment where a white, spatially-coincident auxiliary field is employed and sharpened sensitivities again were found. In both cases, these experimentally-determined sharpened sensitivities agree with our theoretical results. Like Sperling, Foster [23] verified that sharpened sensitivities are a linear combination of the cone sensitivities.

Krastel[24] has measured spectral field sensitivities under changing illumination where, like Sperling, a white conditioning field is employed. Illumination color was changed by placing colored filters in front of the eye. The same test spectral sensitivity curve is measured under both a reddish and bluish illuminant. This suggests that, as in our theoretical experiments, the eye's sharpened mechanisms are unaffected by illumination. More recently Kalloniatis[7] has measured cone spectral sensitivities under white adapting fields of different intensities. The shape of the sharpened sensitivities are independent of the intensity of the adapting field.

Foster argues that the white-background conditioning field and the spatially-coincident auxiliary field inhibit the non-opponent (brightness-based) visual system. He goes on to argue that spectral sharpening is linked to color opponency. Our results, however, suggest a different reason for spectral sharpening; namely, to optimize the sensors so as to decrease their dependence on illumination and enhance color constancy.

6.2 Spectral Sharpening and Double-Opponent cells

At an early stage of post-retinal computation, area V1 in the visual cortex, *double-opponent* (DO) cells have been identified[25]. These cells have spatially- and chromatically-opponent concentric fields, the center and surround, each fed by two cone types. For example, the $R{+}G{-}/G{+}R{-}$ cell has a center which is excited by long-wavelength light and is inhibited by medium-wavelength light. Its surround has an inverse excitation and inhibition.

Hurlbert[26] analyses the operation of the $R{+}G{-}/G{+}R{-}$ cell as:

$$O = G \star \nabla^2 \log(\frac{R}{G}) = G \star \nabla^2 \log(R) - G \star \nabla^2 \log(G) \tag{16}$$

G is a Gaussian (smoothing) filter; $\star$ denotes convolution; ∇^2 is the laplacian operator and O is its output. Equation (16) calculates the laplacian of $\log(\frac{R}{G})$, $\log(\frac{R}{G}) = \log(R) - \log(G)$, at each point in the visual field. Hurlbert suggests equation (16) remains independent of the illumination and represents hue information.

To see how sensor sharpening affects the constancy of equation 16) consider the first directional

derivative of $log(\frac{R}{G})$ at adjacent points x_1 and x_2 which can be approximated by a simple difference.

$$\log(\frac{R_1}{G_1}) - \log(\frac{R_2}{G_2}) = \log(\frac{R_1/G_1}{R_2/G_2}) = \log(\frac{R_1 G_2}{R_2 G_1}) \tag{17}$$

Clearly the constancy, or illumination independence, of equation (17) is related to the constancy of the ratios $\frac{R_1}{R_2}$ and $\frac{G_2}{G_1}$. The sharpened sensors derived in section 5.3 improve the constancy of color ratios. Improving the constancy of the first derivative of $\log(\frac{R}{G})$ implies the constancy of the laplacian calculation of equation (16) will also improve.

Horn[15] in his lightness algorithm uses a convolution operator similar to the DO cell. Specifically, he calculates $O = G \star \nabla^2 \log(S)$ (where $S \in \{R, G, B\}$). This will be illumination invariant to the extent that a DMT is an appropriate vehicle for color constancy. Since spectral sharpening improves the DMT approximation, Horn's results should be more illumination independent when applied to sharpened sensors.

As evident from equations (14) and (15), the sharpened red and green cones are the result of "opponent" interactions. When one of these sharpened sensors is used as the channel in Horn's single-channel Laplacian operator, the result could be considered an opponent-type operator. We see that equation (16) calculates the difference of two Horn-type operators on unsharpened sensors. Hence we predict that Horn's operator with sharpened sensors will be less affected by error than the $R{+}G{-}/G{+}R{-}$ operator. Hurlbert[26] reports evidence of a Horn-type operator, the lightness neurons, residing in the V4 area of the visual cortex.

7 Applications in Machine Vision

In the preceding sections we have demonstrated that spectral sharpening would improve the color constancy performance of DMT algorithms using the human cones as sensors. We propose that spectral sharpening also has application in machine vision.

7.1 Broadband vs Narrowband Filters

The performance of any color constancy algorithm is connected to the choice of sensors. A system designer will likely choose filters which are fairly narrowband—in terms of color constancy the narrower the the better. However, the it is unwise to choose too narrowband a filter or the sensor may by chance not respond at all to a light source with a spiky spectrum such as a fluorescent bulb.

sensors	*% Squared Area*		
	[400, 470]	[460, 500]	[530, 590]
FILTERS	80.0	31.1	53.3
FILTERS sharpened	89.5	77.5	98.8

Table 6: Percentage of squared area, within sharpened intervals, for Kodak filters

There are commercially available filters (e.g. Kodak Wratten filters #93 and #94) having a bandwidth of less than 80 nanometres. Unfortunately, these filters transmit less than 10% of the light so a machine vision system equipped with these filters would require impractically bright lighting to operate. Even less narrowband Wratten filters #25 and #47B, which are commonly used for color separation, absorb most of the incoming light and as such place strong constraints on the lighting conditions. Perhaps narrowband sensitivities could be better achieved by spectral sharpening sets of broadband filters?

As a test, we selected three broadband filters (Wratten #66, #52 and #38)[27] and derived the sharpened filters shown in Figure 9 by applying sensor-based and model-based sharpening. Table 6 contrasts the squared area of the new and original filters with respect to the sharpening intervals. As in the case of sharpening the cone sensitivities, sharpening the Wratten filters has a significant effect. A set of highly transmissive, broadband filters—linear combinations of which are narrowband—would be a good choice for machine color constancy.

8 Conclusion

Spectral sharpening generates sensors which improve the performance of color constancy theories (Von Kries adaptation, Land's retinex, etc.) employing diagonal matrix transformations. Two sharpening techniques have been presented. Model-based sharpening finds sensors which are optimal with respect to a given set of surface reflectances and illuminants. Sensor-based sharpening finds the most narrowband sensors that can be created as a linear combination of the given sensors. The two methods yield quite similar sensor transformations and in both cases substantially improve the accuracy with which a diagonal matrix can model changes in illumination.

Since spectral sharpening aids color constancy, we advance the hypothesis that sharpening might provide a motivation for the psychophysical finding of sharpening in the human visual system.

9 Acknowledgment

This work has been supported by the Simon Fraser University Center for Systems Science and the Natural Sciences and Engineering Research Council of Canada under grant #4322.

10 Figure Captions

Figure 1. Vos Walraven fundamentals (solid line) are contrasted with the sharpened sensitivities derived via Sensor-based (dotted line) and Model-based (dashed line) sharpening.

Figure 2. Vos Walraven sensor responses for surfaces under D100 are mapped to their corresponding descriptors (under D55). Figure 5 is the cumulative histogram of the normalized fitted distance between mapped responses and actual descriptors.

Figure 3. For each illuminant the Cumulative NFD histograms are calculated for the VW curves (solid lines) and for the sharpened VW curves (dashed lines). The top row of Figure 3, from left to right, displays the cumulative histograms for CIE A, D48, and D65. The second row contains the histograms for D75, D100, and the average of all cumulative histograms.

Figure 4 For each illuminant Model-based sharpening returns a different set of sharpened sensors. Figure 4 shows the range of sharpened curves for the VW cone mechanisms.

Figure 5 Experimentally it appears that CIE A has a significant effect on the sharpening ranges of sensors. Figure 5 shows the sharpening range for the VW curves when CIE A is excluded.

Figure 6. Compares performance, by contrasting cumulative NFD histograms, of the VW (solid lines) vs sharpened VW curves (dashed lines). Unlike Figure 3 data fitted with respect to the sharpened curves is transformed back to the original curves for purposes of comparison. That is the performance of the sharpened sensors is judged relative to the original VW sensors.

Let V be the matrix of VW canonical responses and W be the matrix of responses under some other illuminant. We find the matrix $\mathcal{D}$ which minimises the sum of squared errors of $\mathcal{T}V \approx \mathcal{D}\mathcal{T}W$, where $\mathcal{T}$ is the matrix calculated for Sensor-based sharpening (there is a single $\mathcal{T}$ for all illuminants). We then calculate the NFD distances relative to the unsharpened VW sensors: $\mathcal{T}^{-1}\mathcal{T}V \approx \mathcal{T}^{-1}\mathcal{D}\mathcal{T}W$.

Figure 7. The Cumulative NFD histograms for the sharpened VW sensors (dashed lines), calculated as in Figure 6, are compared with the optimal linear transform performance (solid lines).

Figure 8. Sharpened VW curves (solid lines) are contrasted with sharp but slightly flatter alternatives (dashed lines). The short-wave sensor is common to both sensor sets.

Figure 9. Solid lines are filter sensitivities. Dashed lines are sensor-based sharpened filters. Dotted lines are model-based sharpened filters.

References

[1] J. Beck. *Surface Color Perception.* Cornell University Press, 1972.

[2] G. West and M.H. Brill. Necessary and sufficient conditions for von kries chromatic adaption to give colour constancy. *J. Math. Biol.*, 15:249–258, 1982.

[3] D. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, 5:5–36, 1990.

[4] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas.* Wiley, New York, 2nd edition, 1982.

[5] D.H. Foster. Changes in field spectral sensitivities of red-,green- and blue-sensitive colour mechanisms obtained on small backround fields. *Vision Research*, 21:1433–1455, 1981.

[6] H.G. Sperling and R.S. Harwerth. Red-green cone interactions in the increment-threshold spectral sensitivity of primates. *Science*, 172:180–184, 1971.

[7] M. Kalloniatis and R.S. Harwerth. Spectral sensitivity and adaptation characteristics of cone mechanisms under white-light adaptation. *J. Opt. Soc. Am. A*, 7:1912–1928, 1990.

[8] D.C. Hood and M.A. Finkelstein. A case for the revision of textbook models of colour vision. In J.D. Mollon and L.T. Sharpe, editors, *Colour Vision : Physiology and Psychophysics*, pages 385–398. Academic Press, 1983.

[9] J.K. Bowmaker and H.J.A.Dartnell. Visual pigments of rods and cones in the human retina. *J. Physiol.*, 298:501–511, 1980.

[10] D.B. Judd, D.L. MacAdam, and G. Wyszecki. Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.*, 54:1031–1040, August 1964.

[11] R.S. Berns and K.H. Petersen. Empirical modelling of systematic spectrophotometric errors. *Color Res. Appl.*, (4):243, 1988.

[12] C.S. McCamy, H. Marcus, and J.G. Davidson. A color-rendition chart. *J. App. Photog. Eng.*, pages 95–99, 1976.

[13] E.L. Krinov. Spectral reflectance properties of natural formations. *Technical Translation TT-439, National Research Council of Canada*, 1947.

[14] E.H. Land and J.J. McCann. Lightness and retinex theory. *J. Opt. Soc. Amer.*, 61:1–11, 1971.

[15] B. K. P. Horn. Determining lightness from an image. *Computer Vision, Graphics, and Image Processing*, 3:277–299, 1974.

[16] A. Blake. Boundary conditions for lightness computation in Mondrian world. *Computer Vision, Graphics, and Image Processing*, 32:314–327, 1985.

[17] E.H. Land. The retinex theory of color vision. *Scientific American*, pages 108–129, 1977.

[18] J.H. Wilkinson. *Algebraic Eigenvalue Problem.* Monographs on numerical analysis. Oxford University Press, 1965.

[19] D.B. Judd. Hue saturation and lightness of surface colors with chromatic illumination. *J. Opt. Soc. Am.*, 30:2–32, 1940.

[20] D.A. Brainard and B.A. Wandell. Analysis of the reinex theory of color vision. *J. Opt. Soc. Am. A*, 36:1651–1661, 1986.

[21] D.H. Foster. Colour vision. *Contemporary Physics*, 25:477–497, 1984.

[22] D.H. Foster R.S. Snelgar. Test and field spectral sensitivities of colour mechanisms obtained on small white backgrounds: action of unitary opponent-colour processes? *Vision Research*, 23:787–797, 1983.

[23] D.H. Foster and R.S. Snelgar. Initial analysis of opponent-colour interactions revealed in sharpened field sensitivities. In J.D. Mollon and L.T. Sharpe, editors, *Colour Vision : Physiology and Psychophysics*, pages 303–312. Academic Press, 1983.

[24] W. Jaeger H. Krastel and S. Braun. An increment-threshold evaluation of mechanisms underlying colour constancy. In J.D. Mollon and L.T. Sharpe, editors, *Colour Vision : Physiology and Psychophysics*, pages 545–552. Academic Press, 1983.

[25] D.H. Hubel and T.N. Wiesel. Receptive fields ad functional architecture of monkey stirate cortex. *Journal of Physiology (London)*, 195:215–243, 1968.

[26] A.C. Hurlbert. *The Computation of Color (PhD Thesis).* MIT Artificial Intelligence Laboratory, 1989.

[27] Eastman Kodak company. *Kodak filters : for scientific and technical uses*, 2nd edition, 1981.

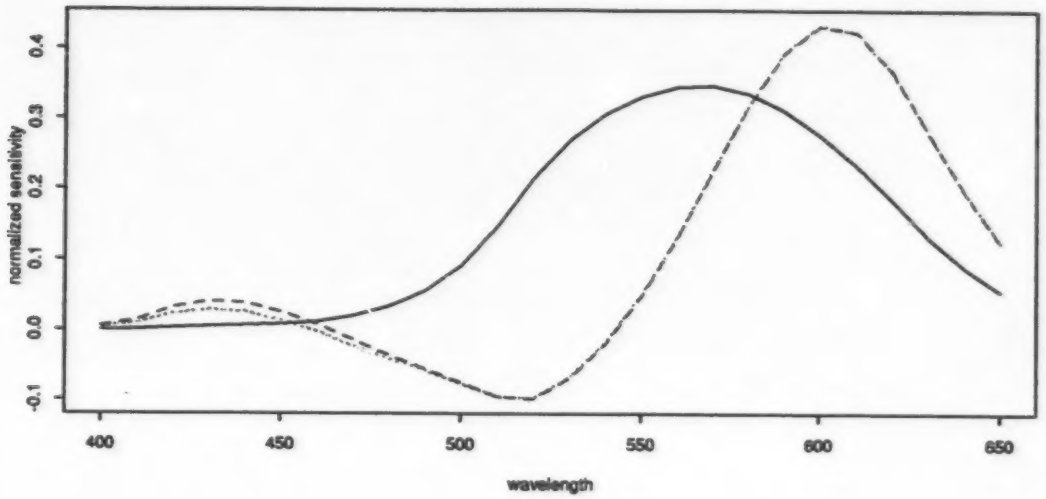
normalized sensitivity
-0.1
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

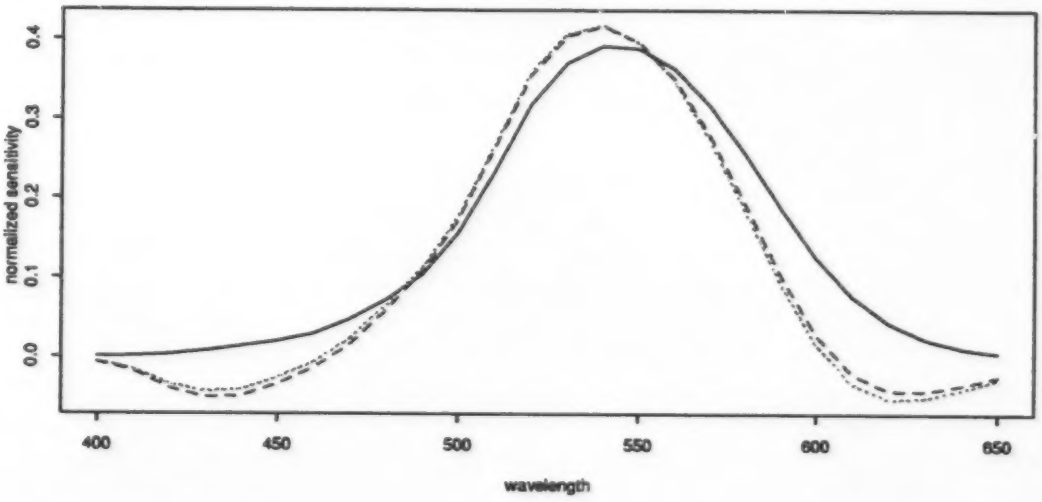
normalized sensitivity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

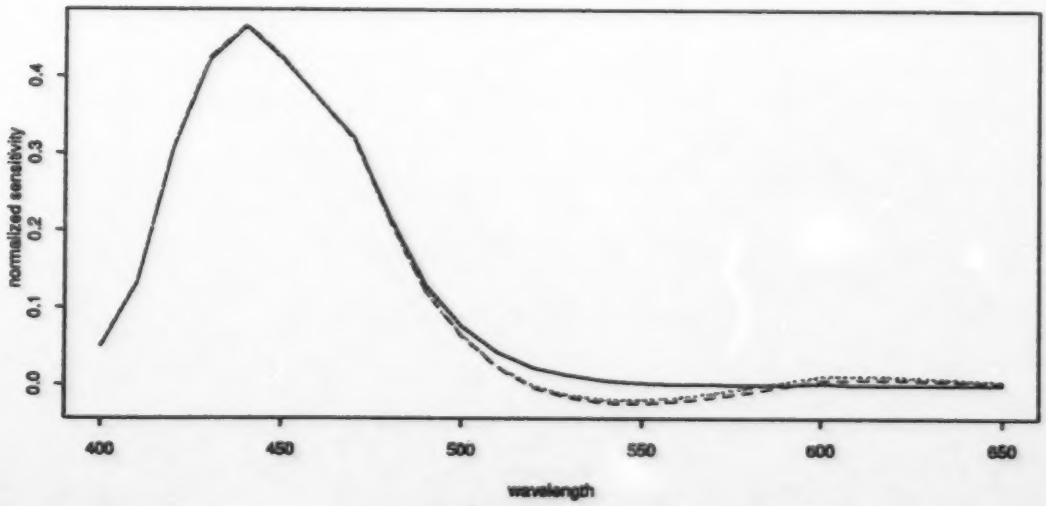
normalized sensitivity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

Figure 1

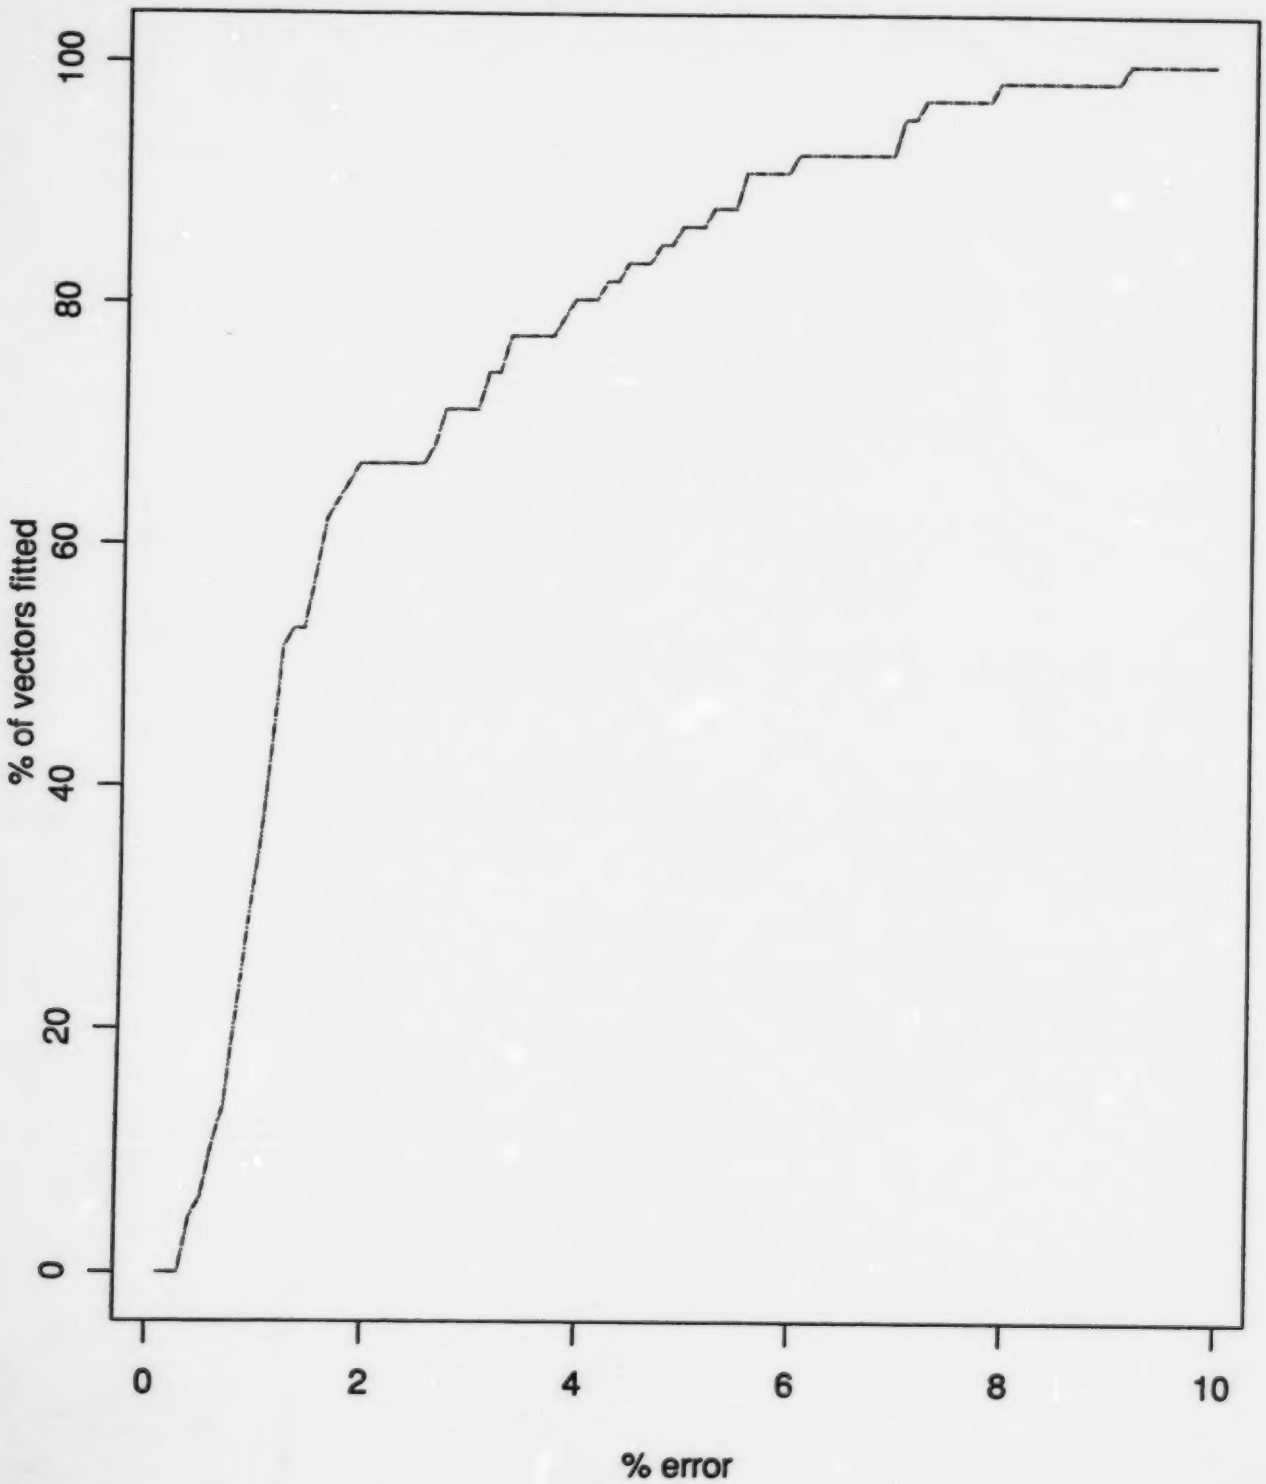

Figure 2

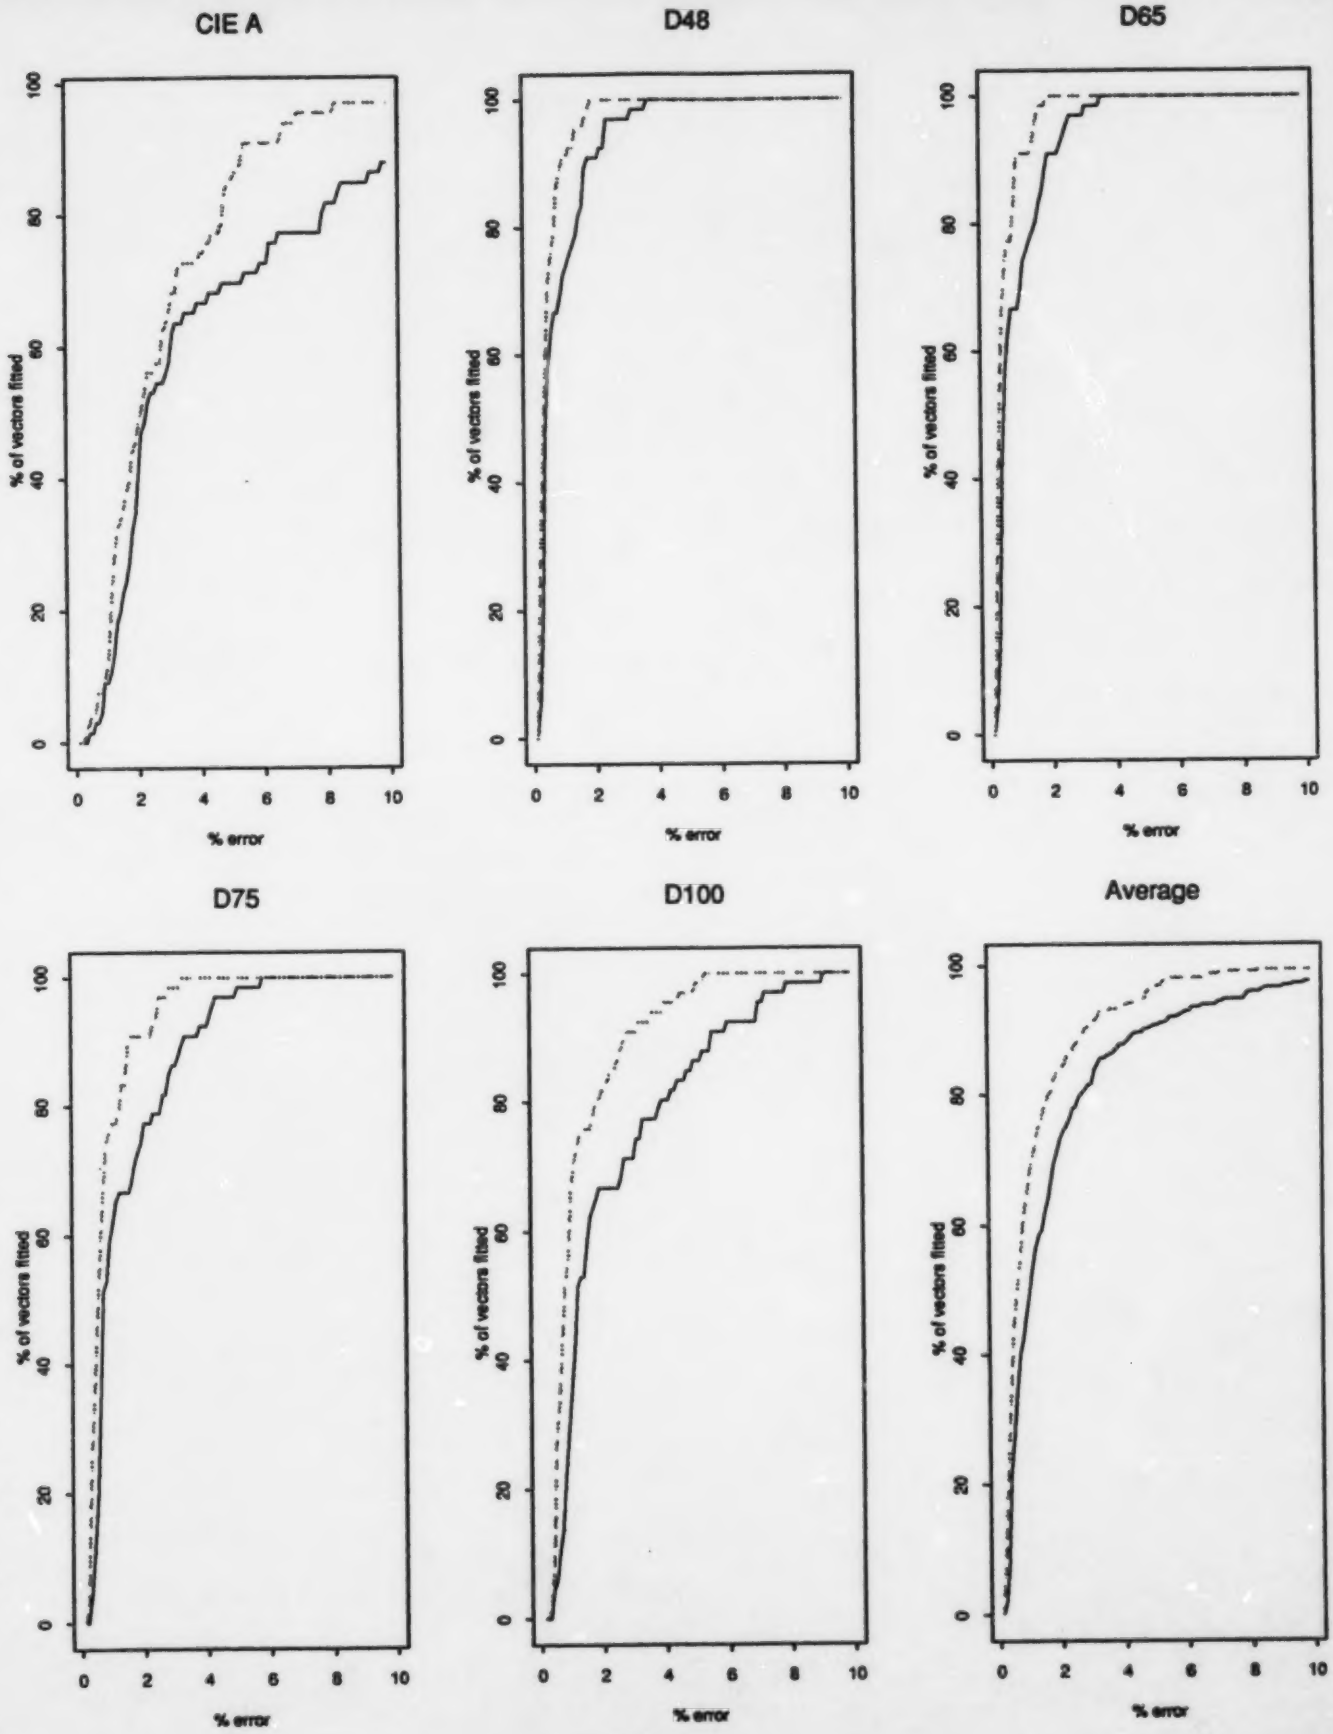
CIE A
D48
D65
D75
D100
Average
% of vectors fitted
% error
0
20
40
60
80
100
0
2
4
6
8
10

Figure 3

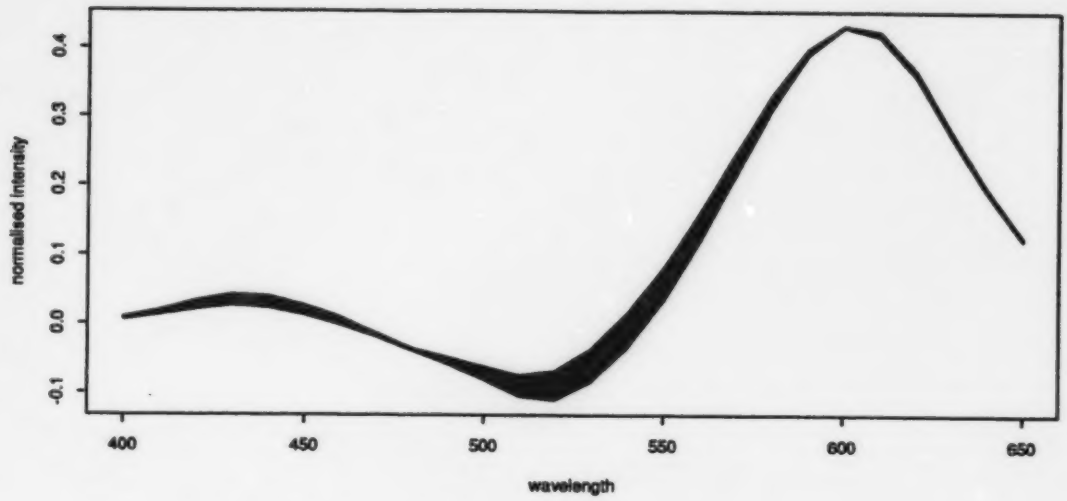
normalised intensity
-0.1
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

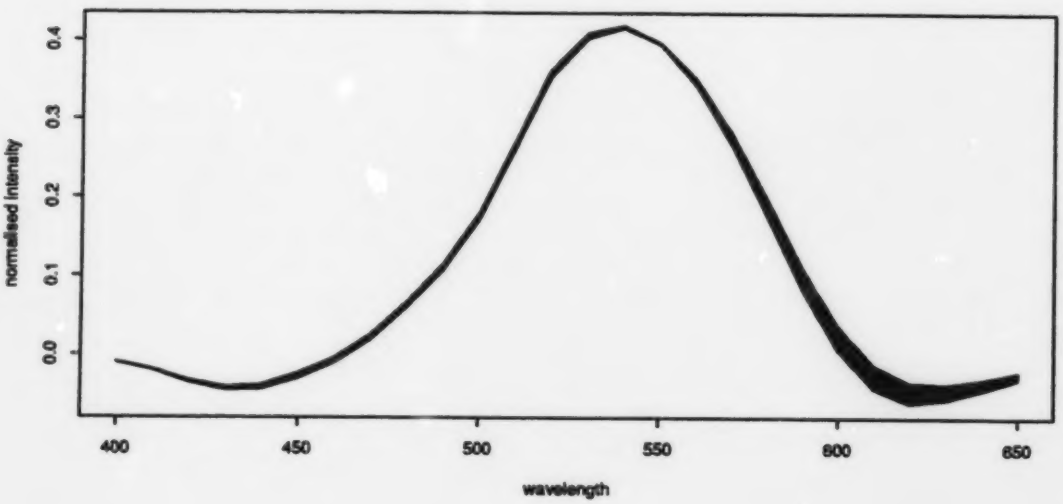
normalised intensity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

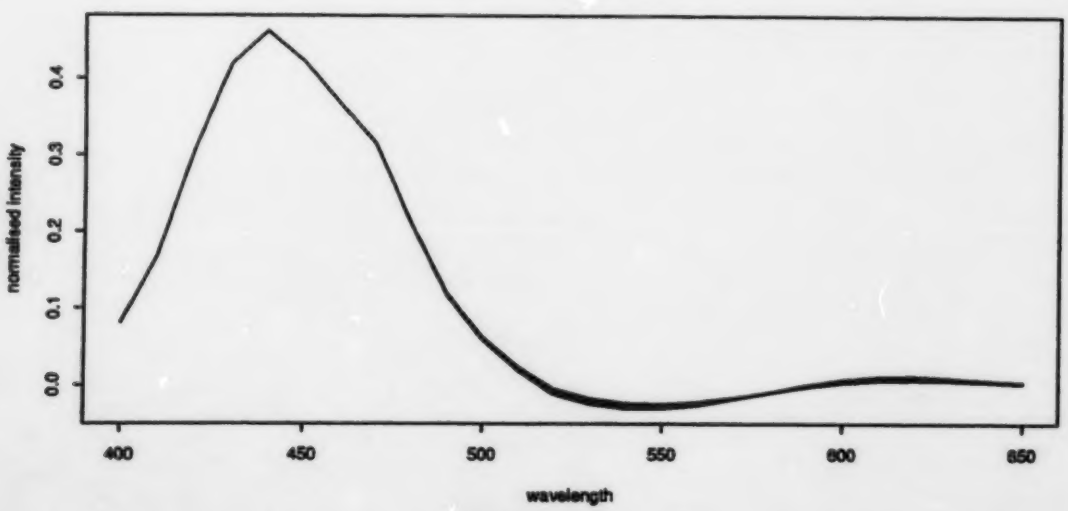
normalised intensity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

Figure 4

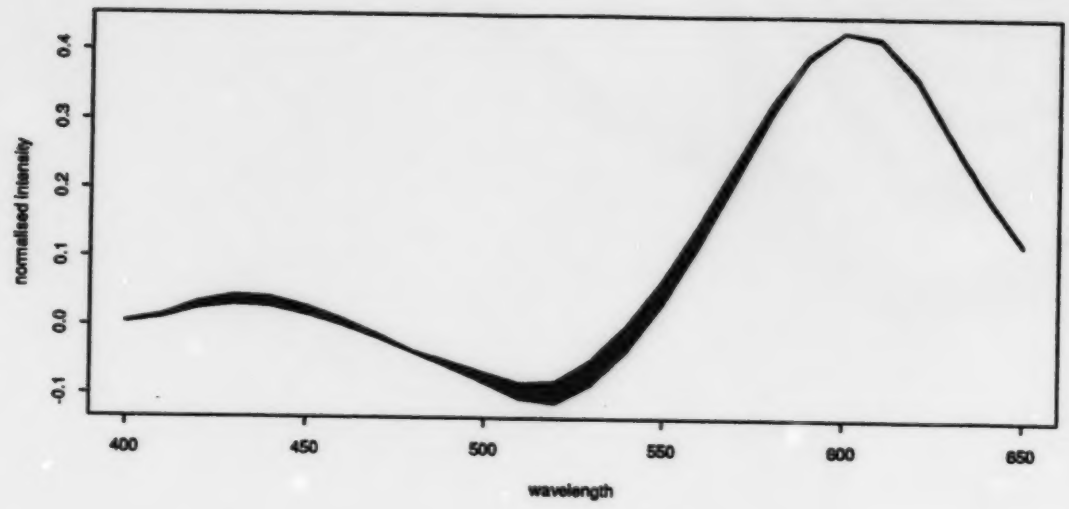
normalised intensity
-0.1
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

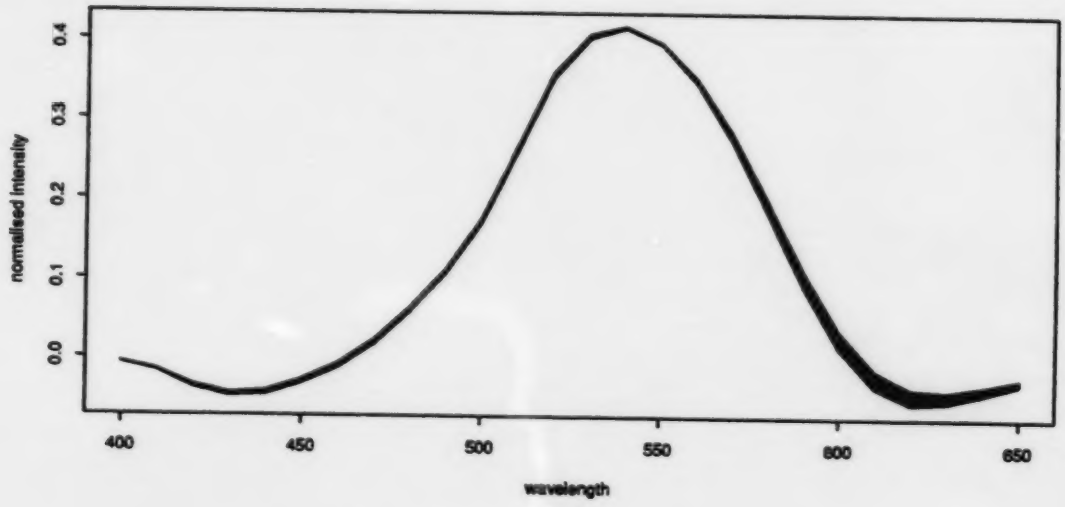
normalised intensity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

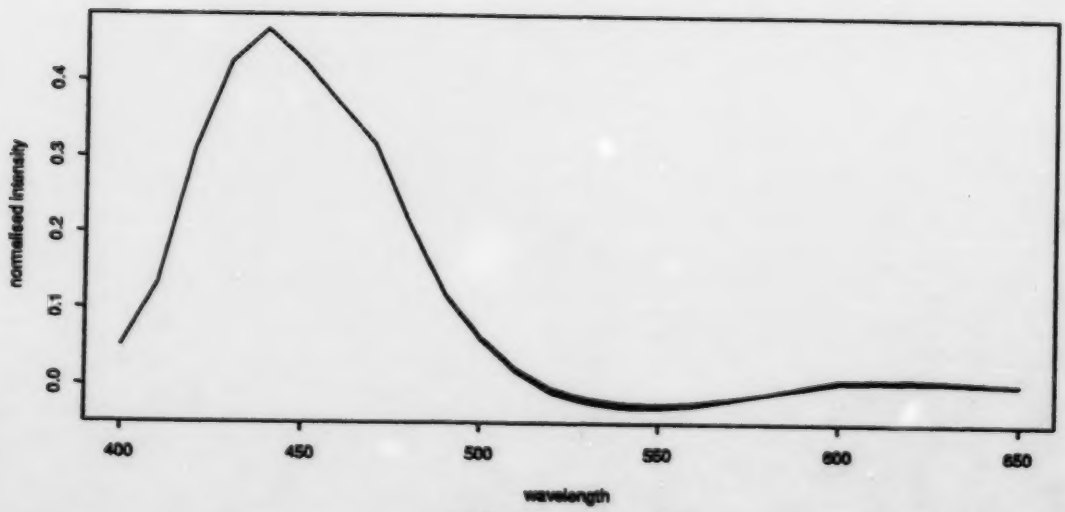
normalised intensity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

Figure 5

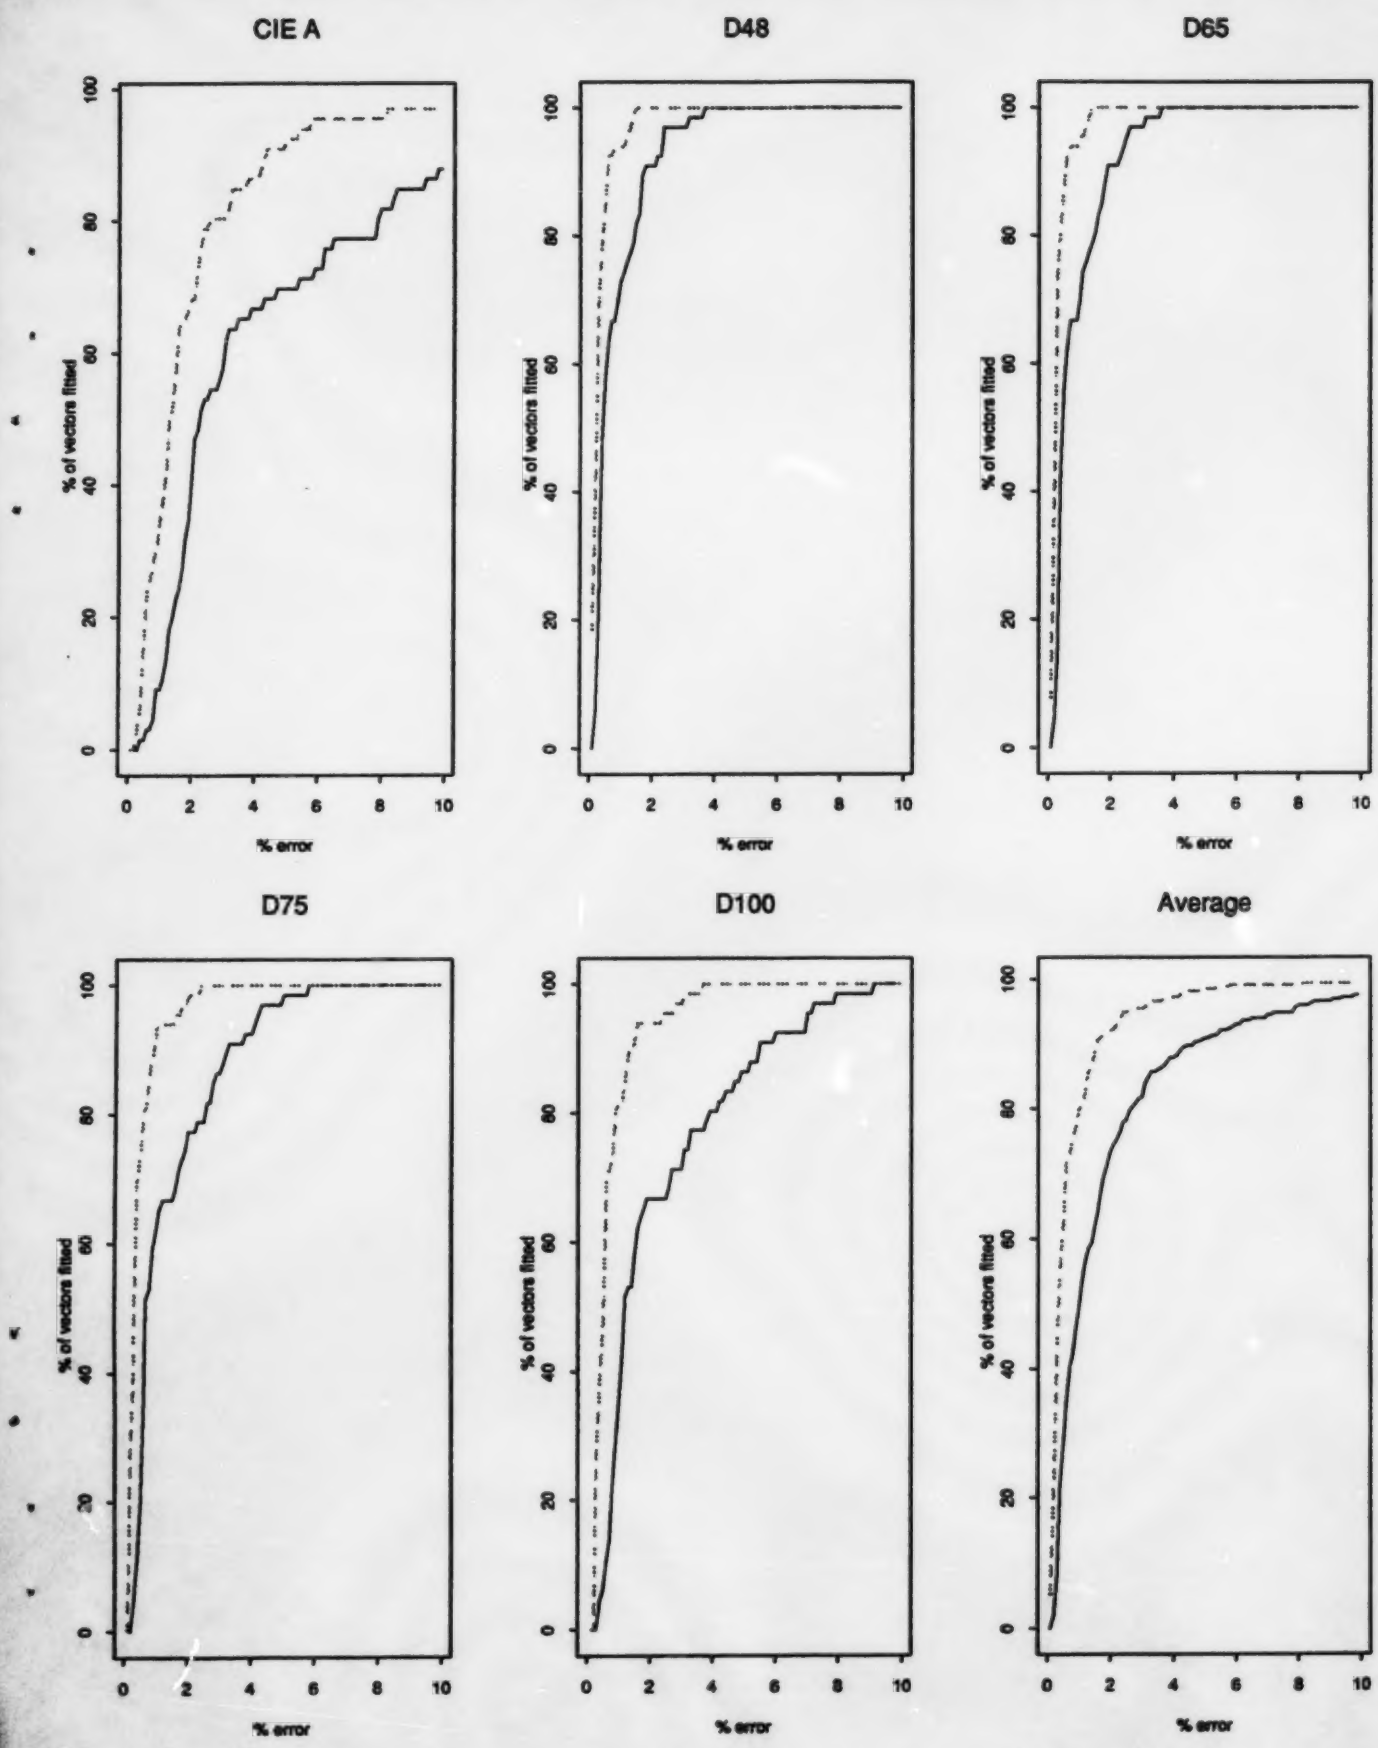
CIE A
D48
D65
D75
D100
Average
% of vectors fitted
% error
0
20
40
60
80
100
0
2
4
6
8
10

Figure 6

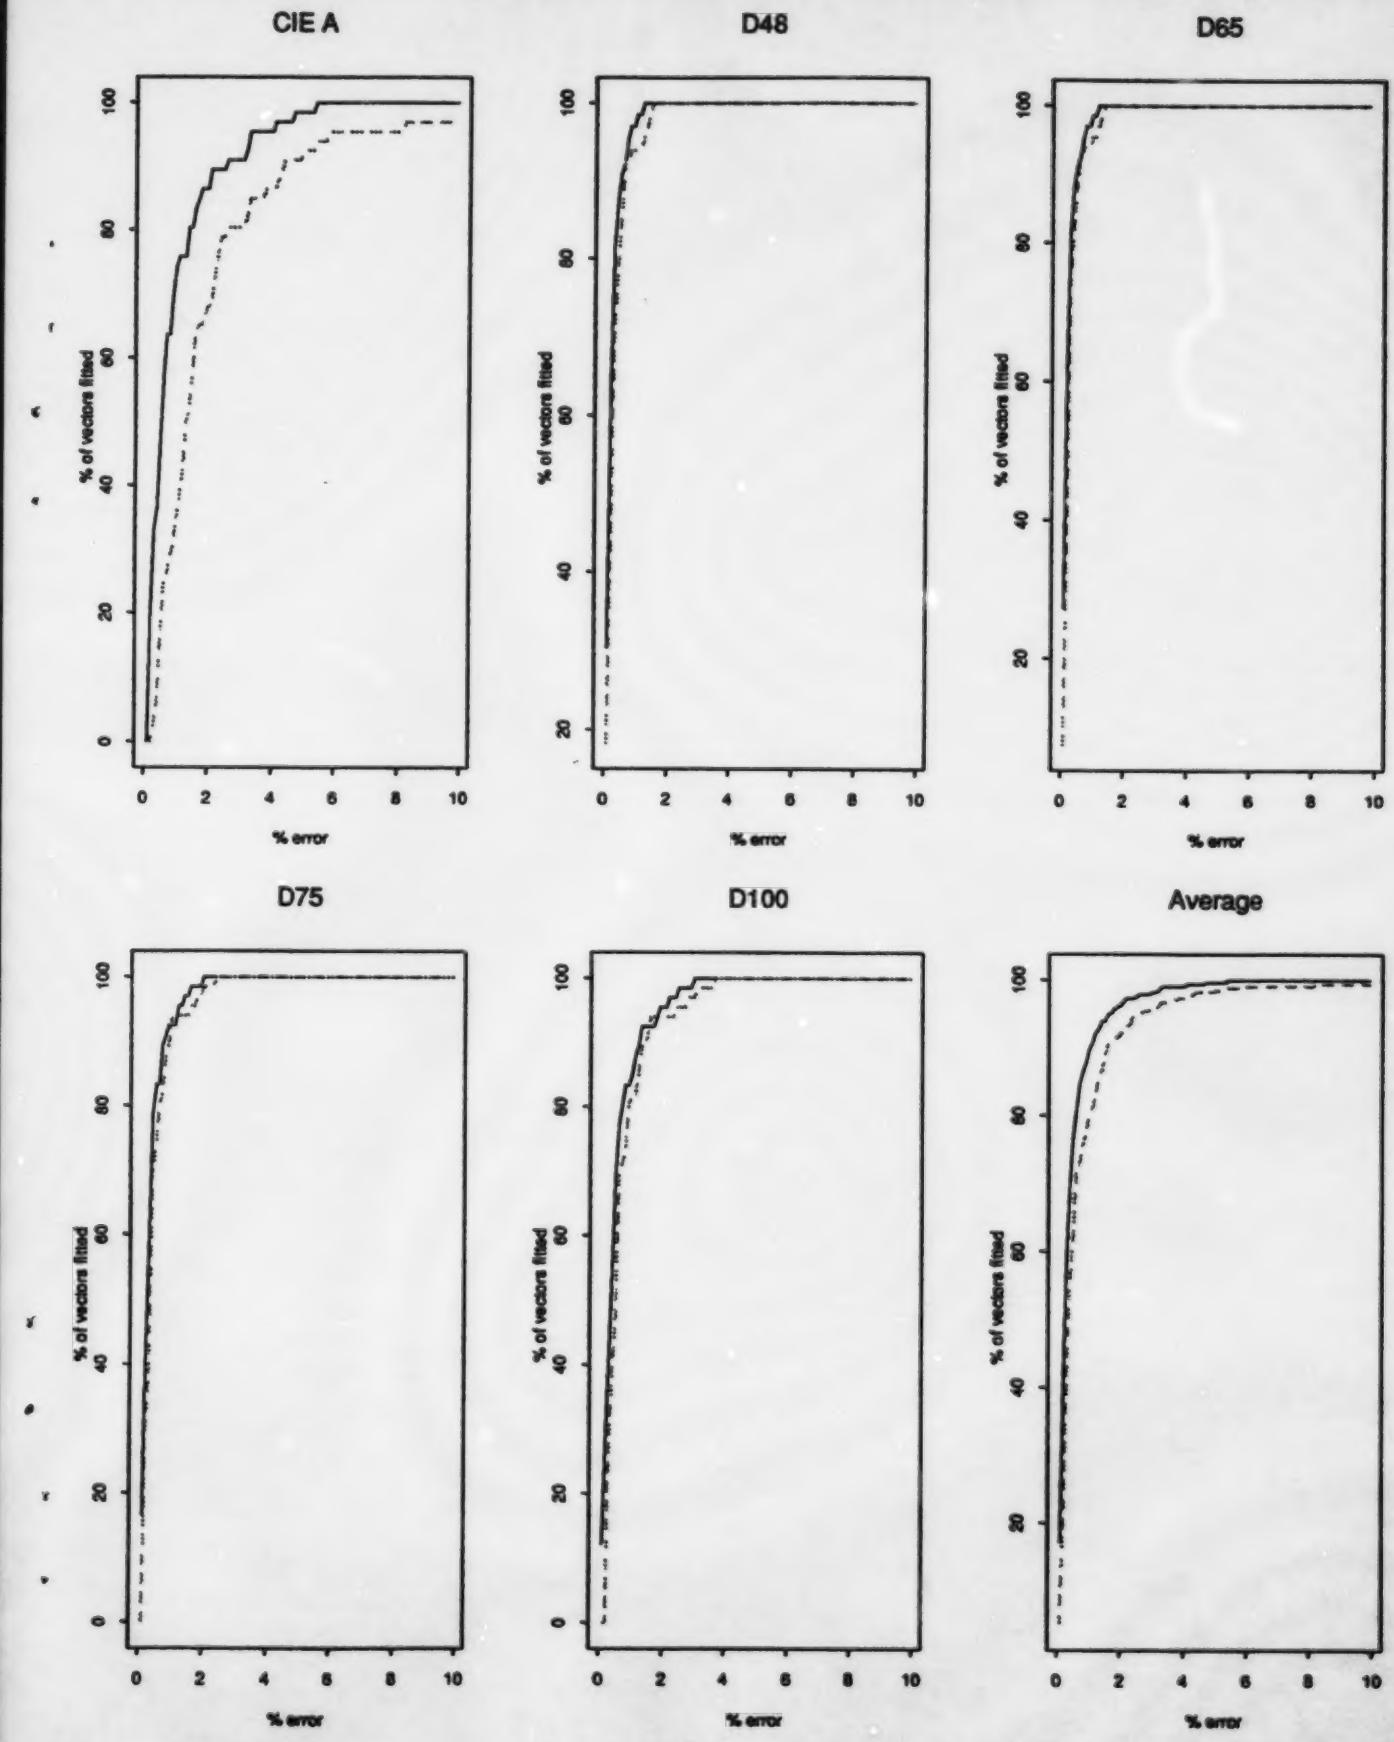
CIE A
D48
D65
D75
D100
Average
% of vectors fitted
% error
0
20
40
60
80
100
0
2
4
6
8
10

Figure 7

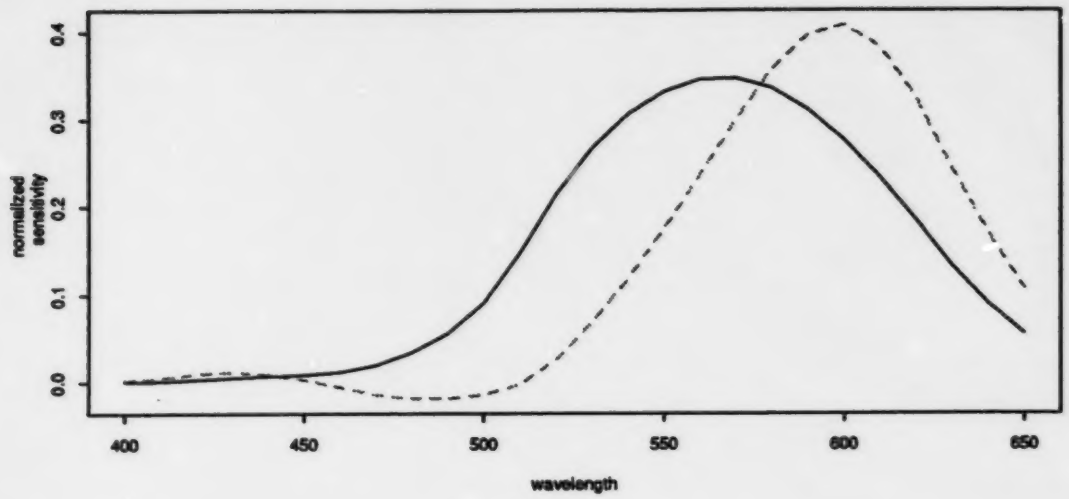
normalized sensitivity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

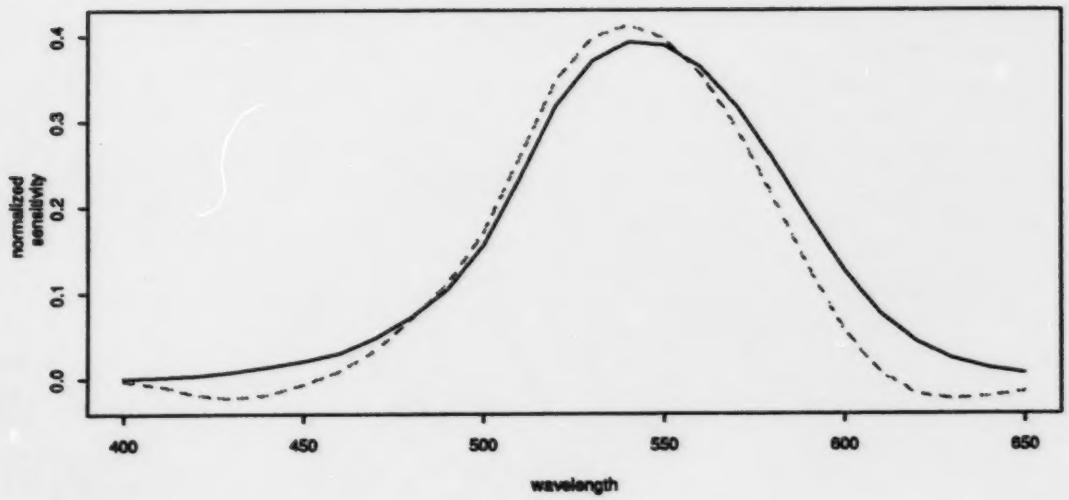
normalized sensitivity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

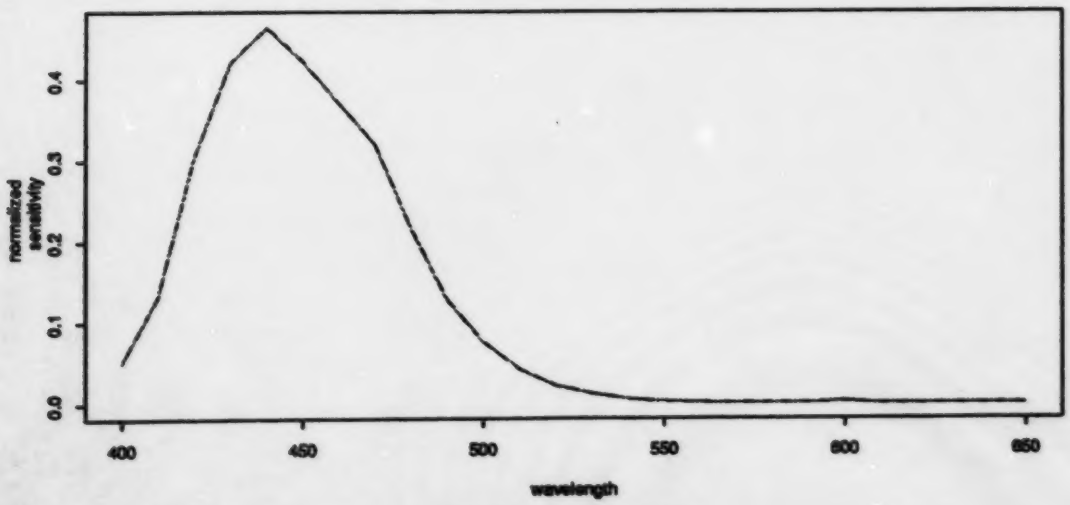
normalized sensitivity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

Figure 8

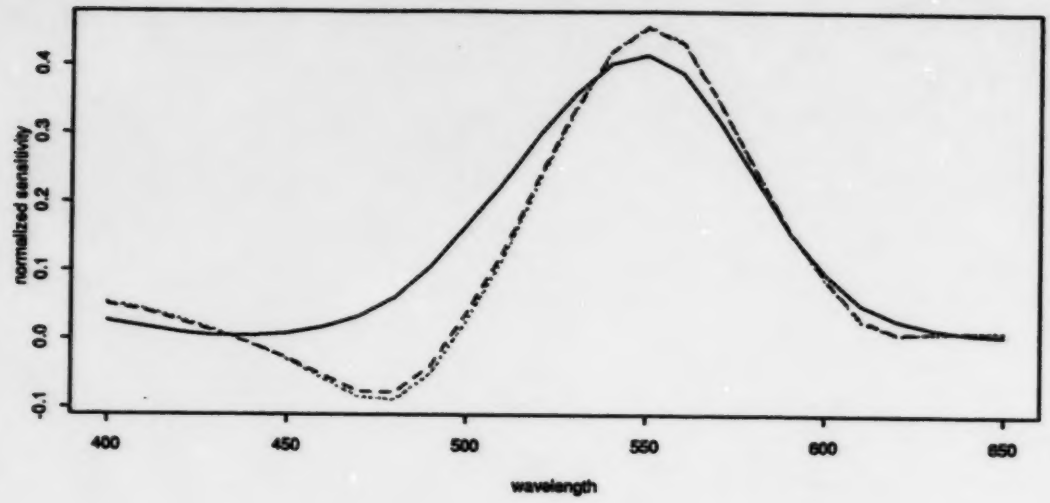
normalized sensitivity
-0.1
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

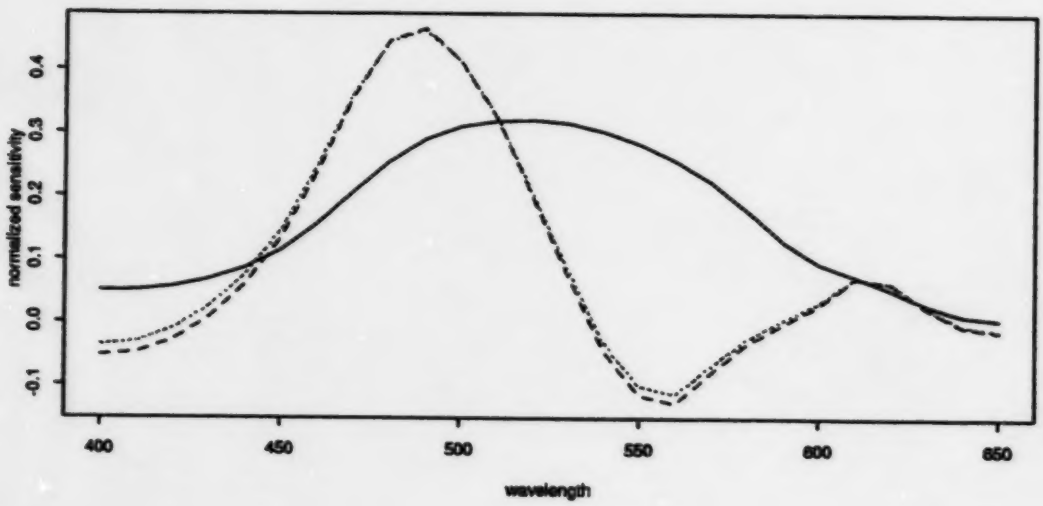
normalized sensitivity
-0.1
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

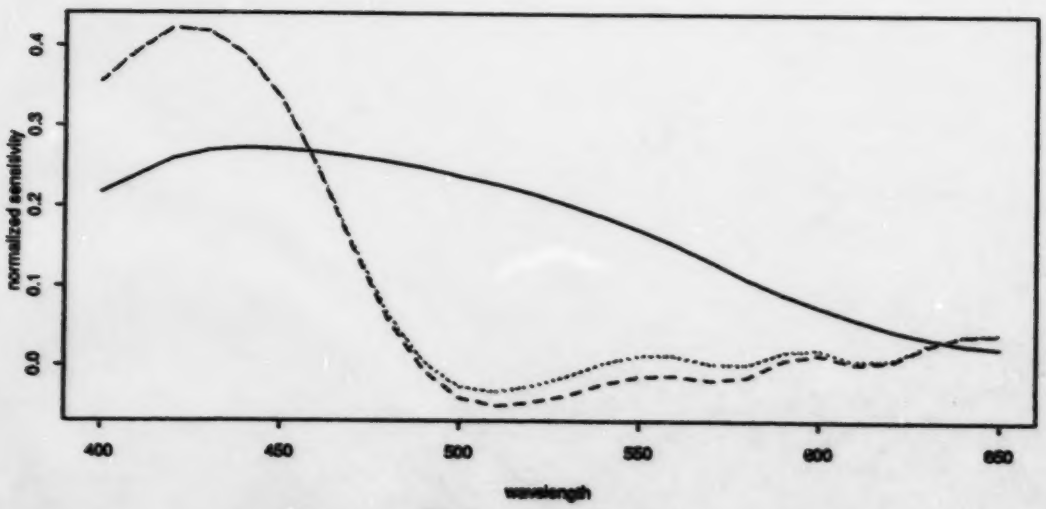
normalized sensitivity
0.0
0.1
0.2
0.3
0.4
400
450
500
550
600
650
wavelength

Figure 9

Centre for Systems Science

• Simon Fraser University established the **Centre for Systems Science** as a provincial centre for advanced study. The objective of CSS is twofold: coordinate and combine advanced systems research at SFU, and encourage application of research results within industry in order to further the technological development of British Columbia.

• CSS has assembled three research groups to streamline its efforts. The orientation of these groups is as follows: intelligent systems, microelectronics, and computer and communications systems. Because of shared research interests, this latter group merges with LCCR (**Laboratory for Computer and Communications Research**) for the publication of technical reports.

• CSS's three groups involve over 80 skilled faculty researchers principally from SFU's Schools of Engineering Science and Computing Science, but also draw substantially on Mathematics, Physics, Geography, Philosophy, Linguistics, Psychology, Communication, and Kinesiology.